

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

**Method And Apparatus For Protectively Operating A
Data/Information Processing Device**

Inventor(s): **Swain W. Porter**

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN, LLP
12400 Wilshire Boulevard, 7th Floor
Los Angeles, California 90025
(503) 684-6200

Method and Apparatus For Protectively Operating A Data/Information Processing Device

BACKGROUND OF THE INVENTION

5

1. **Field of the Invention**

The present invention relates to the field of electronic data/information processing. More specifically, the present invention relates to methods and apparatuses for protectively operating data/information processing devices.

10

2. **Background Information**

The term "data/information processing devices" as used herein is intended to include all microprocessor based devices and/or systems, operated under the control of an operating system. Examples of these devices/systems include but are not limited to general as well as special purpose computing devices/systems, regardless of form factors, palm sized, laptops, desktops, rack mounted, and the like. Examples of special purpose computing devices include but are not limited to set-top boxes, wireless communication devices, and the like. The term "operating system" as used herein is intended to include all software provided to manage and facilitate application usage of hardware resources, however minimal the control and resource scope may be. Typical resource management functions of an "operating system" include task scheduling, memory management and the like. The term "task" as used herein is intended to include its common meaning of an executing instance of a program (a collection of programming instructions).

15

20

25

004814.P014

non-essential programs, such as application programs. The IBM 360 systems provided a supervisor mode and a user mode to segregate privileged system programs and unprivileged user programs. The Multics (Multiplexed Information and Computing Service) developed by Massachusetts Institute of Technology, in 5 cooperation with others, employed a 64 ring approach, combining access node and a triple of ring numbers (r1, r2, r3). In U.S. Patent 4,177,510, issued to Appell et al., a hardware facilitated 4 ring approach is disclosed. Today, the Intel Architecture processors are known to provide a 4 ring hardware facilitated protection through the employment of memory segment descriptors and current task privilege level (CPL). 10 However, partly because most of the other microprocessors remain having a two mode protection approach, the Windows® operating system, used in most Intel Architecture compatible processors, merely employ two of the four ring protection provided by the hardware. The kernel, virtual memory manager and various virtual device drivers (VxD) are executed in ring 0 (the most privileged level), while all other 15 programs, including system services and so forth are executed out of ring 3 (the least privileged level). Rings 1 and 2 are not used.

20 The two levels of protection were reasonably adequate in the days when few programs are executed on most computer systems. Moreover, most of the computer systems operate by themselves, with few interactions from the outside world.

25 Advances in microprocessor, telecommunication and networking technology have dramatically expanded the applications of computing devices, and changed their operating environment. Today, most data/information processing systems are connected to private and/or public networks, such as the Internet, executing programs that are dynamically downloaded from a number of sources. Some

sources are trustworthy, and their programs tend to be well behaved, but others are not.

Accordingly, a need exists to improve the protection of data/information processing systems, especially those operating with a two privilege level protection

5 scheme.

However, this need cannot be easily met, even in the case of systems using Intel Architecture processors and Windows operating system, where there are two unused privileged levels, as the system services and other trustworthy applications are confined to run at the least privileged level (ring 3). It would undermine the

10 stability of the systems, as opposed to increasing its protection, if untrustworthy applications are confined to execute out of the more privileged ring 1 or ring 2.

Relocating the operating system services and other trustworthy programs off the least privileged level (Ring 3) without hardware assistance would require major redesign of the operating system, and raises serious backward compatibility issues.

15 Extending the hardware to have the processor support more privilege levels beyond 4 rings would require major redesign of the processor, as greater than 4 rings would require at least one extra bit be added to the current 2-bit representation. This would cause major redesign to the entire privilege level mechanism, including control register layouts, width of internal data lines, size of comparison circuitry and

20 the like.

Thus, it is further desirable if the need can be met without requiring major processor and/or operating system re-design.

SUMMARY OF THE INVENTION

A privilege level re-mapping mechanism is provided to a processor to re-map privilege levels. The re-mapping mechanism is placed in between the control

5 registers and the privilege checking circuitry, to enable the re-mapping to be dynamically performed in real time prior to privilege checking. The novel dynamic re-mapping of privilege levels prior to privilege checking enables tasks to be executed with relative privilege level relationships that are different from what were nominally assigned to the tasks.

10 In one embodiment, complementary selection mechanism is also provided to enable the novel dynamic re-mapping to be conditionally performed.

ପାତ୍ରାବ୍ଦୀ ପାତ୍ରାବ୍ଦୀ

BRIEF DESCRIPTION OF DRAWINGS

The present invention will be described by way of exemplary embodiments, but not limitations, illustrated in the accompanying drawings in which like references 5 denote similar elements, and in which:

Figure 1 illustrates an overview of the present invention, in accordance with one embodiment;

Figures 2a-2b illustrate the privilege level re-mapper in further detail, in accordance with two embodiments;

10 **Figures 3a-3b** illustrate the privilege level re-mapper in further detail, in accordance with another two embodiments;

Figure 4 illustrates another overview of the present invention, in accordance with another embodiment;

Figure 5 illustrates an example application of the present invention; and

15 **Figure 6** illustrates an example system incorporated with the processor of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

In the following description, various aspects of the present invention will be described. However, it will be apparent to those skilled in the art that the present invention may be practiced with only some or all aspects of the present invention. For purposes of explanation, specific numbers, materials and configurations are set forth in order to provide a thorough understanding of the present invention. However, it will also be apparent to one skilled in the art that the present invention may be practiced without the specific details. In other instances, well known features are omitted or simplified in order not to obscure the present invention.

Parts of the description will be presented using terms such as privilege levels, control registers, and so forth, commonly employed by those skilled in the art to convey the substance of their work to others skilled in the art. Parts of the description will be presented in terms of operations performed by a computer system, using terms such as privilege checks, and so forth. As well understood by those skilled in the art, these quantities and operations take the form of electrical, magnetic, or optical signals capable of being stored, transferred, combined, and otherwise manipulated through mechanical and electrical components of a digital system; and the term digital system include general purpose as well as special purpose data processing machines, systems, and the like, that are standalone, adjunct or embedded.

Various operations will be described as multiple discrete steps performed in turn in a manner that is most helpful in understanding the present invention, however, the order of description should not be construed as to imply that these operations are necessarily order dependent, in particular, the order the steps are presented.

Furthermore, the phrase "in one embodiment" will be used repeatedly, however the phrase does not necessarily refer to the same embodiment, although it may.

Referring now to **Figure 1**, wherein an overview of the present invention in accordance with one embodiment is shown. As illustrated, in accordance with the present invention, a task current privilege level (CPL) remapper **104** that re-maps a task's CPL from one assigned level to another is provided to processor **100**. Task CPL remapper **104** is strategically placed in between task register **102** (a control register where task CPL is stored) and privilege level checking mechanism **106** to enable the re-mapping to be dynamically performed in real time prior to privilege checking during execution. The novel dynamic re-mapping of privilege levels prior to privilege checking advantageously enables tasks to be executed with relative privilege level relationships that are different from what were nominally assigned to the tasks.

Except for the teachings of the present invention incorporated, processor **100** is otherwise intended to represent a broad range of processors known in the art. As will be readily apparent from the descriptions to follow, while **Fig. 1** specifically illustrates the task register where a task's current privilege level, the present invention applies to privilege level in general, and may be practiced to dynamically alter the relative privilege relationship between memory segments, selectors, descriptors and the like. For ease of understanding, the remaining description will nevertheless continue to primarily refer to the task's CPL. Privilege checking mechanism **106** is intended to represent a broad range of privilege checking mechanisms or circuitry known in the art. It may enforce any one of a number of privilege rules, as well as enforcing these rules in any one of a number of implementation manner. Neither the privilege rules being enforced nor the manner

they are enforced are of particular relevance to the practice of the present invention. In fact, a major advantage of the present invention is the ability to introduce a new order of privilege relationship without requiring major re-design to the fundamental privilege protection mechanism of a processor nor the operating system that uses
5 the processor.

Figures 2a-2b illustrate task CPL re-mapper **104** in further detail, in accordance with two embodiments. **Fig. 2a** illustrates a basic embodiment, where a new control register **202** is used to re-map a task's CPL. As illustrated, the re-
10 targeted privilege levels are stored in register **202**, and they are selectively accessed and retrieved using the task CPL read out of task register **102** as an offset into register **202**. As a result, a task having a CPL of "0" or "1" will retain the "0" or "1" CPL, whereas a task with a CPL of "2" will be re-mapped to a CPL of "3", and a task with a CPL of "3" will be re-mapped to a CPL of "2". Accordingly, the desired
15 privilege level re-mapping, and relative privilege relationship re-ordering is achieved in accordance with the stored scheme.

Fig. 2b illustrates a more elaborate embodiment, where a memory storage array **204** is used to re-map a task's CPL. As illustrated, multiple sets of re-targeted privilege levels are stored in array **204**, and they are selectively accessed and
20 retrieved using the task CPL read out of task register **102** as a row pointer into array **204**, in conjunction with a configuration signal serving as a column pointer into array **204**. As a result, a task having a CPL of "0", "1", "2" or "3" may be re-mapped to "0", "1", "3" and "2" respectively as before, if the set stored in column 1 is used, or to "1", "0", "3" and "2" respectively, if the set stored in column 4 is used instead.
25 Accordingly, the desired privilege level re-mapping, and relative privilege relationship re-ordering is achieved in accordance with one of the stored schemes.

The re-targeted privilege levels representing a re-mapping scheme may be "hard coded" into register **202** or array **204**, or it may be loaded at power-on or reset as part of the initialization process. The configuration signal may be driven e.g. off a programmable configuration register (not shown).

5 Thus, it can be seen from the embodiments of **Fig. 2a-2b**, the present invention may be practiced with a simple pre-determined re-mapping scheme or with a re-mapping scheme to be configurally determined from a rich or full set of all possible re-mappings.

10 **Figures 3a-3b** illustrate task CPL re-mapper **104** in further detail, in accordance with another two embodiments. **Fig. 3a** achieves the same re-mapping as the embodiment of **Fig. 2b** if the set of re-targeted privilege levels stored in column 4 are used. Except, under **Fig. 3a**, the re-mapping is achieved through a combinatorial circuit element, XOR gate **302**. More specifically, the lower order bit 15 of a task's CPL is XOR'd with the value "0" to alter it, while the higher order bit is retained. Effectively, a task with CPL "00", "01", "10" and "11" (0, 1, 2, 3 in decimal) will be re-mapped to "01", "00", "11" and "10" (1, 0, 3, 2 in decimal).

20 **Fig. 3b** may achieve a number of re-mappings possible under the earlier described embodiments. Except, under **Fig. 3b**, the re-mapping is also achieved through a combinatorial circuit element, XOR gate **302**. In addition to XOR gate **302**, the embodiment of **Fig. 3b** is also provided with selector **304** to allow the selective retaining of the original lower bit or the employment of the altered lower bit. Selector **304** selects either the original lower bit or the altered lower bit in accordance with a configuration signal. The original lower bit (or the altered lower 25 bit) may be selected with the configuration signal equals "0" or "1". The manner of selection is immaterial. Thus, if configuration signal always selects the original

lower bit, re-mapping is effectively disabled. On the other hand, if configuration signal is conditionally driven to select the altered lower bit, depending on whether the lower bit is "1" or "0", it achieves the same re-mapping offered by the embodiment of **Fig. 2b** employing column 1 (which is the same as the embodiment of **Fig. 2a**), or the re-mapping offered by the embodiment of **Fig. 2b** employing column 3. Finally, if configuration signal always selects the original altered lower bit, the embodiment of **Fig. 3b** is effectively the same as the embodiment of **Fig. 3a**.

Similarly, configuration signal may be driven from a programmable configuration register, or outputs of other combinatorial circuits. Thus, it can be seen that various re-mapping may also be achieved through combinatorial circuits. The embodiments of **Fig. 3a-3b** are kept simple for ease of understanding. However, those skilled in the art will be able to extend from these embodiments to allow even more flexible re-mapping of various kinds.

Figure 4 illustrates another overview of the present invention, in accordance with another embodiment. The embodiment of **Fig. 4** is essentially that of the embodiment of **Fig. 1**, except for the provision of selector **402** to allow the re-mapping to be selectively enabled and disabled. In other words, the inclusion of selector **402** enables the present invention to be configurably included or excluded.

Figure 5 illustrates an example application of the present invention. As illustrated, in this example application, the kernel, the virtual device driver and the memory manager of an operating system are nominally attributed with task CPL "0", enabling them to execute in privilege ring 0, whereas other operating system services, as well as "trustworthy" applications are nominally attributed with task CPL "3", confining them to execute in privilege ring 3. Untrustworthy applications, such

as Internet applications, are nominally attributed with task CPL "2", enabling them to execute in the more privileged ring 2.

However, employing the present invention, the privilege levels are dynamically re-mapped, enabling the relocation of the operating system services 5 and trustworthy applications to the more privileged ring 2, and confining the untrustworthy Internet application to the least privileged ring 3 instead.

What constitutes trustworthiness is application dependent. Their demarcation is immaterial for the practice of the present invention. Further, the term "privilege 10 ring" or "ring" as used herein is intended to include its conventional meaning that a program afforded a more inner privilege ring typically has privileges inclusive that of another program afforded a more outer privilege ring.

Thus, it can be seen under the present invention, a class of lesser privileged tasks can be carved out of the existing least privileged tasks. The new least privileged tasks will first be nominally given a more privileged level. But, at 15 execution time, the privilege levels of the residual former least privileged tasks and the new least privileged tasks are re-mapped (prior to privilege checking), and re-ordered to the desired relative privilege relationship. Likewise, the same may be performed at the other end of the privilege spectrum. That is, a class of more privileged tasks can be carved out of the existing most privileged tasks. The new 20 more privileged tasks will first be nominally given a lesser privilege level. But, at execution time, the privilege levels of the residual former most privileged tasks and the new more privileged tasks are re-mapped (prior to privilege checking), and re-ordered to the desired relative privilege relationship.

25 Referring now to **Figure 6**, wherein a block diagram illustrating an example system incorporated with the teachings of the present invention is shown. System

600 is intended to represent a broad range of digital systems or devices known in the art, including but not limited to computer systems of all form factors (from palm-sized, to laptop, desktop and racked mounted servers), telecommunication devices such as wireline or wireless telephones, or entertainment devices such as set top devices, and the like. As shown, example system **600** includes processor **602**, system memory **604** coupled to each other via "bus" **612**. Coupled also to "bus" **612** are non-volatile storage **606**, input/output device **608** and communication interface **610**.

Processor **602** may be the processor of **Fig. 1**, **Fig. 4**, and other equivalents.

10 Each of the other enumerated elements is intended to represent a wide range of the respective devices/elements known in the art. For example, system memory **604** may be SDRAM, DRAM and the like, from semiconductor manufacturers such as Micron Technology of Boise, Idaho. Bus **612** may be a single bus or a multiple bus implementation. In other words, bus **612** may include multiple buses of identical or different kinds properly bridged, such as Local Bus, VESA, ISA, EISA, PCI and the like. Non-volatile storage **606** may be disk drives or CDROMs from manufacturers such as Seagate Technology of Santa Cruz of CA, and the like. Input/Output devices **608** may include input devices, such as keypads, key boards, or cursor control devices like a mouse, a track ball and so forth, from vendors such as

20 Logictech of Milpitas, CA, and output devices like display devices such as LCD displays, flat panel displays or monitors of any types, from vendors such as Viewsonic of Walnut, CA. Communication interface **610** may be a wireless interface, or a wireline interface, such as modem interface, an ISDN adapter, a DSL interface, an Ethernet or Token ring network interface and the like, from vendors

25 such as 3COM of San Jose, CA.

Thus, a method and apparatus for protectively operating a data/information processing system has been described. While the present invention has been described in terms of the above illustrated embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments described. The present invention can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of restrictive on the present invention.

Wildseed - M&A For Protectively
Operating A Data/Information ...